UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/991,219	11/16/2001	Manabu Kitamura	16869S-038000US	3936

20350 7590 09/02/2005

TOWNSEND AND TOWNSEND AND CREW, LLP
TWO EMBARCADERO CENTER
EIGHTH FLOOR
SAN FRANCISCO, CA 94111-3834

EXAMINER	
GILLIS, BRIAN J	

ART UNIT	PAPER NUMBER
2141	

DATE MAILED: 09/02/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/991,219	KITAMURA ET AL.	
	Examiner	Art Unit	
	Brian J. Gillis	2141	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *13 June 2005*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *19-40* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *19-40* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *16 November 2001* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other: _____.

DETAILED ACTION

Continued Examination Under 37 CFR 1.114

A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on June 13, 2005 has been entered.

Claim Objections

Claim 40 is objected to because of the following informalities: The claim refers to the storage server of claim 34, which claims 37, an exact copy of this claim already depends from. The Examiner interprets this as a typographical error and should be dependent on claim 39. For examination purpose the Examiner assumes dependency on claim 39. Please notify the Examiner if this is incorrect. Appropriate correction is required.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 19-40 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ofek et al (US Patent #5,680,640) in view of Khattar et al (NPL).

Claims 19, 27, 28 disclose a storage server for providing a virtualized storage apparatus to a plurality of hosts, the storage server comprising: a first communication interface coupled to a first network to receive a data access request from one or more hosts (figure 1, ref. 28, column 2, lines 18-25); a second communication interface (figure 1, ref. 20a and 20b) coupled to a second network to communicate with first (figure 1, ref. 14) and second storage (figure 1, ref. 16) subsystems including first and second storage areas, respectively, to store data associated with the hosts (figure 1, ref. 14, 16, and 17a-n and column 4, lines 13-17; wherein the devices are in a network environment, hence, each device is connected to each other via communication interfaces), the first and second storage subsystems being provided at remote locations from the storage server (p. 3, paragraph 3); and a virtual device driver component that is operable to present a virtual storage area to the hosts, the virtual storage area being mapped (figure 1, ref. 24) to the first and second storage areas of the first and second storage subsystems (column 6, lines 18-36; wherein the virtual device driver component is the component that makes the virtual storage devices perform, completely transparent to the hosts, a "background" data migration process using the data map to check data have been migrated and which is still to be migrated) wherein the storage server is configured to receive a first data request including a virtual address identifying a first location in the virtual storage area from one of the hosts (column 4, lines 13-17; wherein the devices are in a network environment, hence, each device has a unique ID recognized by the other devices) and generate a second data request directed to the first storage subsystem, the second data request including a storage address identifying

a second location in the first storage area that is mapped to the first location in the virtual storage area (abstract; when the host requests data from a data storage, the system determines if the data requested is stored on the first or second storage area. If the data is not on the new storage system or second storage system, the process will look for the data on the old storage system or first storage system. Once migrated the data are mapped to the corresponding place so the user can access it transparently without knowing on which storage area to look for the data), wherein the storage server is configured to change mapping of the first location in the virtual storage area to a third location in the second storage area (column 6, lines 37-45; wherein the map/table is configure d to change mapping of the old storage location to the new storage location) in connection with a data migration operation (column 2, lines 40-55 and column 6, lines 17-25; wherein the data access read/write are handled while the migration is running as the "background" process). Ofek et al discloses the use of direct fiber optic connections to connect the devices. It fails to disclose a switch for interconnecting the host computers, storage apparatuses, and the back end server. Ofek et al also fails to disclose the first and second storage subsystems being provided at remote locations from the storage server. Khattar et al teaches the use of switches to connect a large number of network devices (p. 15, section 1.3.3.12) including storage devices. Khattar et al also teaches of the storage subsystems being provided at remote locations from the storage server (p. 3, paragraph 3).

Ofek et al and Khattar et al are analogous art because they are both related data migration.

At the time of the invention it would have been obvious to a person of ordinary skill in the art to use the switches and storage subsystems in Khattar et al with the system in Ofek et al because storage area networks allow applications that move data to perform better (Khattar et al p. 4, paragraph 2).

Claim 20 discloses the storage server of claim 19, wherein the data migration involves migration of data in the second location to the third location. Ofek et al further teaches wherein the data are migrated from the old storage to the new storage with the updated map to indicate the old data has been migrated (abstract).

Claim 21 discloses the storage server of claim 20, wherein if the storage server receives a third data request including the virtual address identifying the first location of the virtual storage area from one of the hosts after the data migration, the storage server is configured to generate a fourth data request directed to the second storage subsystem, the fourth data request including a storage address identifying the third location in the second storage area. Ofek et al further teaches wherein the data is migrated from the old storage to the new storage with the updated map to indicate the old data has been migrated. The data map is keeping track of the migration and the first, second, and third location in the new storage area with the address. The map must have an addressing mechanism in order to keep track of where the data is migrated (abstract and figure 2).

Claim 22 discloses the storage server of claim 19, wherein the first and second network switches are the same. Khattar et al further teaches the switches allow any to any connections (page 15-16, section 1.3.3.12).

Claim 23 discloses the storage server of claim 19, wherein the server is configure d to initiate the data migration of first data stored in the first storage area of the first storage subsystem to the second storage area of the second storage subsystem. Ofek et al further teaches if the server is coordinating the migration and keeping a data map of the migration state, it must initiate the migration (abstract).

Claim 24 discloses the storage server of claim 23, wherein the server is configured to receive a request from one of the hosts to access a portion of the first data while the first data are being migrated to the second storage area Ofek et al shows a request is received while data is migrated (figure3, ref. 104, 106, 108).

Claim 25 discloses the storage server of claim 24, wherein the server accesses a data migration unit to determine whether to direct a server request generated in response to the request from one of the hosts to access the first data is to be directed to the first or second storage subsystem, wherein the data migration unit maintains a record on portions of the first data that remain in the first storage area as the first data are being migrated to the second storage area. Ofek et al further teaches the server determines if the data is in the old or new storage and direct the access accordingly (column 6, lines 17-34 and figure 1, ref. 24 [Migrate process] and figure3, ref. 102).

Claim 26 discloses the storage server of claim 25, wherein the data migration unit is provided in one of the storage subsystem or the server. Ofek et al further teaches the migration unit is in the server (figure1, ref. 24).

Claim 29 discloses the method of claim 28, further comprising: maintaining a table holding a flag for indicating a data migration state on each data block associated

with the virtualized storage area. Ofek et al further teaches a data element map including an indication of whether or not an element is stored on a storage system (column 2, lines 28-31 and figure2).

Claim 30 discloses the method claim 28, wherein each data block corresponds to a logical block address. Ofek et al further teaches the system uses fixed block architecture (column 4, lines 54-59).

Claim 31 discloses the method claim 28, wherein the table is maintained by one of the first and second storage subsystems. Ofek et al further teaches the second storage device includes a data element map (column 2, lines 28-32).

Claim 32 discloses the method claim 31, further comprising: receiving a third data request at the backend server including the virtual address identifying the first location of the virtual storage area from a second host after the data migration; and generating a fourth data request directed to the second storage subsystem, the fourth data request including a storage address identifying the third location in the second storage area. Ofek et al further teaches the data is migrated from the old storage to the new storage with the updated map to indicate the old data has been migrated. The data map is keeping track of the migration and the first, second, and third location in the new storage area with the address. The map must have an addressing mechanism in order to keep track of there the data are migrated to (abstract and figure 2).

Claim 33 discloses the method of claim 32, wherein the first and second hosts are different hosts. Ofek et al further teaches the first and second hosts are different (column 5, lines 19-20).

Claim 34 discloses a storage server for providing a virtualized storage apparatus to a plurality of hosts, the storage server comprising: a first communication interface coupled to a first network to receive a data access request from one or more hosts that are coupled to the first switch (figure 1, ref. 28 and column 2, lines 18-25); a second communication interface (figure 1, ref. 20a and 20b) coupled to a second network to communicate with first (figure 1, ref. 14) and second storage (figure 1, ref. 16) subsystems including first and second storage areas, respectively, to store data associated with the hosts (figure 1, ref. 14, 16, and 17a-n and column 4, lines 13-17; wherein the devices are in a network environment, hence, each device is connected to each other via communication interfaces), the first and second storage subsystems being provided at remote locations from the storage server (p. 3, paragraph 3); and a virtual device driver component that is operable to present a virtual storage area to the hosts, the virtual storage area being mapped (figure 1, ref. 24) to the first and second storage areas of the first and second storage subsystems (column 6, lines 18-36; wherein the virtual device driver component is the component that makes the virtual storage devices perform, completely transparent to the hosts, a "background" data migration process using the data map to check data have been migrated and which is still to be migrated) wherein the storage server is configure d to receive a first data request including a virtual address identifying a first location in the virtual storage area from one of the hosts via the first communication interface (column 4, lines 13-17; wherein the devices are in a network environment, hence, each device has a unique ID recognized by the other devices) and generate a second data request directed to the

first storage subsystem, the second data request including a storage address identifying a second location in the first storage area that is mapped to the first location in the virtual storage area (abstract; when the host requests data from a data storage, the system determines if the data requested is stored on the first or second storage area. If the data is not on the new storage system or second storage system, the process will look for the data on the old storage system or first storage system. Once migrated the data are mapped to the corresponding place so the user can access it transparently without knowing on which storage area to look for the data), wherein the storage server is configure d to change mapping of the first location in the virtual storage area to a third location in the second storage area (column 6, lines 37-45; wherein the map/table is configure d to change mapping of the old storage location to the new storage location) in connection with a data migration operation (column 2, lines 40-55 and column 6, lines 17-25; wherein the data access read/write are handled while the migration is running as the "background" process), and wherein the first and second network switches are the same. Ofek et al teaches of the use of direct fiber optic connections to connect the devices. It fails to disclose a switch for interconnecting the host computers, storage apparatuses, and the back end server. Ofek et al also fails to disclose the first and second storage subsystems being provided at remote locations from the storage server. Khattar et al teaches the use of switches to connect a large number of network devices (p. 15, section 1.3.3.12) including storage devices. Khattar et al also teaches of the storage subsystems being provided at remote locations from the storage server (p. 3, paragraph 3).

Ofek et al and Khattar et al are analogous art because they are both related data migration.

At the time of the invention it would have been obvious to a person of ordinary skill in the art to use the switches and storage subsystems in Khattar et al with the system in Ofek et al because storage area networks allow applications that move data to perform better (Khattar et al p. 4, paragraph 2).

Claim 35 discloses the storage server of claim 34, wherein the data migration involves migration of data in the second location to the third location. Ofek et al further teaches wherein the data are migrated from the old storage to the new storage with the updated map to indicate the old data has been migrated (abstract).

Claim 36 discloses the storage server of claim 35, wherein if the storage server receives a third data request including the virtual address identifying the first location of the virtual storage area from one of the hosts after the data migration, the storage server is configure d to generate a fourth data request directed to the second storage subsystem, the fourth data request including a storage address identifying the third location in the second storage area. Ofek et al further teaches wherein the data is migrated from the old storage to the new storage with the updated map to indicate the old data has been migrated. The data map is keeping track of the migration and the first, second, and third location in the new storage area with the address. The map must have an addressing mechanism in order to keep track of where the data is migrated (abstract and figure2).

Claim 37 discloses the storage server of claim 34, wherein the server is configure d to initiate the data migration of first data stored in the first storage area of the first storage subsystem to the second storage area of the second storage subsystem. Ofek et al further teaches if the server is coordinating the migration and keeping a data map of the migration state, it must initiate the migration (abstract).

Claim 38 discloses the storage server of claim 37, wherein the server is configure d to receive a request from one of the hosts to access a portion of the first data while the first data are being migrated to the second storage area Ofek et al shows a request is received while data is migrated (figure3, ref. 104, 106, 108).

Claim 39 discloses a storage server for providing a virtualized storage apparatus to a plurality of hosts, the storage server comprising: a first communication interface coupled to a network switch to receive a data access request from one or more hosts that are coupled to the switch (figure 1, ref. 28 and column 2, lines 18-25); a second communication interface (figure 1, ref. 20a and 20b) coupled to the network switch to communicate with first (figure 1, ref. 14) and second storage (figure 1, ref. 16) subsystems including first and second storage areas, respectively, to store data associated with the hosts (figure 1, ref. 14, 16, and 17a-n and column 4, lines 13-17; wherein the devices are in a network environment, hence, each device is connected to each other via communication interfaces), the first and second storage subsystems being provided at remote locations from the storage server (p. 3, paragraph 3); and a virtual device driver component that is operable to present a virtual storage area to the hosts, the virtual storage area being mapped (figure 1, ref. 24) to the first and second

storage areas of the first and second storage subsystems (column 6, lines 18-36; wherein the virtual device driver component is the component that makes the virtual storage devices perform, completely transparent to the hosts, a "background" data migration process using the data map to check data have been migrated and which is still to be migrated) wherein the storage server is configure d to receive a first data request including a virtual address identifying a first location in the virtual storage area from one of the hosts via the first communication interface (column 4, lines 13-17; wherein the devices are in a network environment, hence, each device has a unique ID recognized by the other devices) and generate a second data request directed to the first storage subsystem, the second data request including a storage address identifying a second location in the first storage area that is mapped to the first location in the virtual storage area (abstract; when the host requests data from a data storage, the system determines if the data requested is stored on the first or second storage area. If the data is not on the new storage system or second storage system, the process will look for the data on the old storage system or first storage system. Once migrated the data are mapped to the corresponding place so the user can access it transparently without knowing on which storage area to look for the data), wherein the storage server is configure d to change mapping of the first location in the virtual storage area to a third location in the second storage area (column 6, lines 37-45; wherein the map/table is configure d to change mapping of the old storage location to the new storage location) in connection with a data migration operation (column 2, lines 40-55 and column 6, lines 17-25; wherein the data access read/write are handled while the migration is running as

the "background" process), wherein the data migration involves migration of data stored in the second to the third location (abstract), wherein if the storage server receives a third request including the virtual address identifying the first location of the virtual storage area from one of the hosts via the first interface after the data migration, the storage server is configure d to generate a fourth data request directed to the second storage subsystem and send the fourth data request to the second storage subsystem via the second interface, the fourth data request including a storage address identifying the third location in the second storage area (abstract and figure2). Ofek et al teaches of the use of direct fiber optic connections to connect the devices. It fails to disclose a switch for interconnecting the host computers, storage apparatuses, and the back end server. Ofek also fails to disclose the first and second storage subsystems being provided at remote locations from the storage server. Khattar et al teaches the use of switches to connect a large number of network devices (p. 15, section 1.3.3.12) including storage devices. Khattar also teaches of the storage subsystems being provided at remote locations from the storage server (p. 3, paragraph 3).

Ofek et al and Khattar et al are analogous art because they are both related data migration.

At the time of the invention it would have been obvious to a person of ordinary skill in the art to use the switches and storage subsystems in Khattar et al with the system in Ofek et al because storage area networks allow applications that move data to perform better (Khattar et al p. 4, paragraph 2).

Claim 40 discloses the storage server of claim 39, wherein the server is configure d to initiate the data migration of first data stored in the first storage area of the first storage subsystem to the second storage area of the second storage subsystem. Ofek et al further teaches if the server is coordinating the migration and keeping a data map of the migration state, it must initiate the migration (abstract).

Response to Arguments

Applicant's arguments with respect to claims 19, 27, and 28 have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Fiacco et al (US Patent #6,098,125) teaches of mapping fibre channel frames based on control and type header fields. Yanai et al (US Patent #5,544,347) teaches of a data storage system controlled remote data mirroring with respectively maintained data indices.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Brian J. Gillis whose telephone number is 571-272-7952. The examiner can normally be reached on M-F 7:30-5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Rupal Dharia can be reached on 571-272-3880. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Brian J Gillis
Examiner
Art Unit 2141

BJG

RUPAL DHARIA
SUPERVISORY PATENT EXAMINER